REMARKS

In the Office Action dated June 14, 2005, claims 61, 62, 64 and 65 are rejected under 35 U.S.C. 102(e) as being anticipated by Yeldener (US Patent No. 5,999,897).

As described in pages 40 and 41 of the specification of the present patent application, the over-sampled synthesis signal $\hat{S}$ (Figure 2) does not contain the higher frequency components which were lost by the down-sampling process (module 101 of Figure 1) at the encoder 100. This gives a low-pass perception to the synthesized speech signal. To restore the full band of the original signal, a high frequency generation procedure is disclosed. This procedure is performed in modules 201 to 216, and adder 221, and requires input from voicing factor generator 204 (Figure 2).

Claim 61 of the present patent application broadly defines this high frequency generation procedure as follows:

> 61. A device for recovering a high frequency content of a wideband signal previously down-sampled and for injecting said high frequency content in an over-sampled synthesized version of said wideband signal to produce a full-spectrum synthesized wideband signal, said high-frequency content recovering device comprising:
>
> a) a random noise generator for producing a noise sequence having a given spectrum;
>
> b) a spectral shaping unit for shaping the spectrum of the noise sequence in relation to linear prediction filter coefficients related to said down-sampled wideband signal; and
>
> c) a signal injection circuit for injecting said spectrally-shaped noise sequence in the over-sampled synthesized signal version to thereby produce said full-spectrum synthesized wideband signal.

In the Office Action, the Examiner indicates that Yeldener discloses (in Fig. 1 (7, 8, 9, 12, 13) and Fig. 2B) a device for recovering a high frequency content of a wideband signal previously down-sampled and for injecting the high frequency content in an over-sampled synthesized version of the wideband signal to produce a full-spectrum synthesized wideband signal.

The following comments evidence that Yeldener does not describe a high frequency content recovering device but, as described in column 1, lines 44-47, a device for estimating pitch of a

speech signal using perception based analysis by synthesis which provides a very robust performance and is independent of the input speech signals (column 1, lines 44-47).

Figure 2B of Yeldener illustrates a noise spectrum generator. In column 4, lines 59-64, Yeldener describes that, for the unvoiced part of the excitation spectrum, a white random noise spectrum normalized to excitation band energies, is used for the frequency components that fall above the cut-off frequency ($\omega>\omega_c$). The voiced and unvoiced excitation signals are then added together to form the overall synthesized excitation signal.

In Figure 2B of Yeldener, the noise spectrum is subjected, in series, to energy normalization, inverse FFT and "overlap & add" operations. Therefore, the noise spectrum is not shaped in relation to linear prediction filter coefficients. The LPC filter of Figure 2B is the synthesis filter of the decoder. This synthesis filter generates the reconstructed speech and, therefore, is not related to shaping the spectrum of a noise sequence in view of producing or recovering a previously lost high frequency signal content.

In Figure 2B of Yeldener, the adder operation adds the voiced and unvoiced excitation signals to form the overall synthesized excitation signal. The resultant excitation is then shaped by a linear time-varying LPC filter to form the final synthesized speech. In order to enhance the output speech quality and make it cleaner, a frequency domain post-filter is used.

Accordingly, Figure 2B of Yeldener fails to describe injection of a spectrally-shaped noise sequence in a synthesized signal version to thereby produce the full-spectrum synthesized wideband signal.

Regarding Figure 1, Yeldener in column 3, lines 21-53, describes an Analysis By Synthesis error minimization procedure which is applied to choose the most optimal pitch estimate. First, a segment of speech signal S(n) is analyzed in an LPC analysis section 3 where linear predicitive coding (LPC) is used to obtain LPC filter coefficients. The segment of speech is then passed through an LPC inverse filter 4 using the estimated LPC filter coefficients in order to provide a residual signal which is spectrally flat.

Contrary to the examiner's suggestion, the LPC inverse filter is not a spectral shaping unit for shaping the spectrum of a noise sequence. The input signal S(n) to the filter 4 is defined by Yeldener as being an input speech signal and not a random noise sequence.

In Figure 1 of Yeldener, the residual signal from filter 4 is multiplied by a window function W(n) at multiplier 5 and transformed into the frequency domain to provide a residual spectrum using either DFT (or FFT) in a DFT section 6. Next, in peak picking section 7, the residual spectrum is analyzed to determine the peak amplitudes and corresponding frequencies and phases. In a sinusoidal synthesis section 8, the peak components are used to generate a reference residual (excitation) signal. The reference residual signal is then passed through an LPC synthesis filter 9 to obtain a reference speech signal. The envelope or spectral shape of the residual spectrum is calculated in a spectral envelope section 10, the envelope of the residual spectrum is sampled at the harmonics of the corresponding pitch candidate to determine the harmonic amplitudes and phases for each pitch candidate in a harmonic sampling section 11, these harmonic components are provided to a sinusoidal synthesis section 12 where they are used to generate a harmonic synthetic residual (excitation) signal, and the synthetic residual signal for each pitch candidate is then passed through a LPC synthesis filter 13 to obtain a synthetic speech signal for each pitch candidate. Each of the synthetic speech signals are then compared with the reference signal in an adder 14 to obtain a signal to noise ratio for each of the synthetic speech signals.

Accordingly, the chain 4-5-6-7-8-9 of Figure 1 of Yeldener is used to generate a reference speech signal. Chain 4-5-6-10-11-12-13 of Figure 1 of Yeldener is used to generate a synthetic speech signal. Therefore, none of these two chains is used, as defined in claim 1 of the present patent application, to produce a high frequency content of a speech signal.

As discussed hereinabove, Figure 1 of Yeldener does not produce a noise sequence having a given spectrum, does not shape the spectrum of the noise sequence in relation to linear prediction filter coefficients and, therefore, cannot inject such a spectrally-shaped noise sequence in a synthesized signal version to produce a full-spectrum synthesized signal.

The adder 14 of Figure 1 of Yeldener calculates a signal to noise ratio. It is not used to inject a spectrally-shaped noise sequence in a synthesized signal version to produce a full-spectrum synthesized wideband signal.

The above comments also apply to all the other independent claims.

The claims have been amended to define, as suggested by the examiner, all the elements in all the equations.

In view of the above, Applicant believes the pending application is in condition for allowance.

Dated: December 13, 2005

Respectfully submitted,

By Laura C Brutman

Laura C. Brutman
Registration No.: 38,395
DARBY & DARBY P.C.
P.O. Box 5257
New York, New York 10150-5257
(212) 527-7700
(212) 527-7701 (Fax)
Attorneys/Agents For Applicant